

# NASA Deep Space Network: Automation Improvements in the Follow-the-Sun Era

Mark D. Johnston, Michael Levesque, Shan Malhotra, Daniel Tran, Rishi Verma, and Silvino Zendejas

Jet Propulsion Laboratory, California Institute of Technology

4800 Oak Grove Drive, Pasadena CA USA 91109

{mark.d.johnston, michael.levesque, shan.malhotra, daniel.tran, rishi.verma, silvino.zendejas} @jpl.nasa.gov

## Abstract

The Deep Space Network (DSN) comprises three sites, located in California, Spain, and Australia; each site operates one 70m and multiple 34m antennas that provide communications and navigation services to highly elliptical and deep space missions. The DSN is operated by JPL for NASA, and serves both US and international missions. As part of a multiyear upgrade in automation of the network, JPL has undertaken a project called “Follow the Sun Operations” (FtSO), which will fundamentally change the operations paradigm of the DSN. In this new operations model, each one of the three sites will operate the entire network during their day shift, handing off control to the next site as their day ends. This is in contrast to the current approach, wherein each site operates only their local antennas and equipment, but does so 24 hours/day, 7 days/week. The FtSO model offers the potential for significant operations cost savings, but poses some unique challenges as operations shifts from local to remote. This paper discusses some of these FtSO challenges in the areas of increased automation related to complexity management, reactive re-scheduling, and improved monitoring and situational awareness.

## 1. Introduction

The NASA Deep Space Network (DSN) consists of three large complexes of antennas, spaced roughly evenly in longitude around the world at Goldstone, California; Madrid, Spain; and Canberra, Australia. Each complex contains one 70 meter antenna along with a number of 34 meter and smaller antennas, as well as the electronics and networking infrastructure to command and control the antennas and to communicate with various mission control centers. Figure 1 shows two of the 34 meter antennas at the Canberra Deep Space Communications Complexes (DSCC), which currently operate four of the 13 antennas in the network, with a fifth under construction for initial use in 2016. For more extensive background on the DSN, refer to [1,2].

All NASA planetary and deep space missions, as well as many international missions, communicate to Earth through the DSN. In some cases, missions closer to Earth also use the DSN, some routinely, others on an occasional basis. The capabilities of the DSN make it a highly capable scientific facility in its own right, so it is used for radio astronomy (including very long baseline interferometry) as well as radio science investigations. At present, there are

45 regular distinct users of DSN, who together schedule about 500 activities per week on 13 antennas. Over the next few decades, utilization of the DSN is expected to grow significantly, with more missions operating, higher data rates and link complexities, and the possibility of manned mission support. In addition, the total number of antennas will grow to 18 by the mid 2020s, while at the same time there is pressure to reduce ongoing costs yet maintain an around-the-clock operational capability.

Presently, each of the DSN complexes is staffed 24x7 with local personnel who manage the antenna/spacecraft links. The individuals directly responsible for this are designated Link Control Operators, or LCOs. In general, each LCO manages up to two links at a time. Future plans for increased automation are presently in progress, under the general term “*Follow the Sun*” Operations (FtSO), which includes the following two fundamental shifts in operational paradigm:

- *Remote Operations (RO)* — at each complex during their local day shift, each complex will operate not only their local assets, but also all the assets of the other two complexes as well, via remote control.
- *Three Links per Operator (3LPO)* — the number of links a LCO will manage will increase from two (today) to three.

These changes represent a major paradigm shift and will require numerous software changes to improve DSN automation, as well as WAN upgrades to increase bandwidth and reliability of complex-to-complex communications. The benefit will be a significant savings in operations costs while continuing to provide high-quality support to DSN users.

In this paper we first give an overview of the DSN automation improvements planned to support the evolution to FtSO (Sect 2). We concentrate on the Service Preparation phase of the system ([3], which addresses planning, scheduling, and support product generation prior to service execution. Sect 3 includes a brief description of the changes to be introduced by the Follow-the-Sun paradigm shift. We then describe our approach to helping the operations staff manage the complexity of remote operations as well as the increasing number of simultaneous links per operator (Sect 4). Finally, we conclude with a discussion of plans for next steps (Sect 5).



Fig. 1 Two of the four 34 meter antennas at the Canberra, Australia, Deep Space Communications Complex. In addition, the facility also includes another 34 meter antenna and also a 70 meter antenna.

## 2. Follow-the-Sun Operations — Overview

There are two major aspects of the FtSO paradigm shift: remote operations, and three links per operator. Fig 2 shows an overview schematic diagram of the remote operations network management and handoff in Follow-the-Sun era. Each complex controls not just their own antennas and associated assets, but all those of the other two complexes as well. Eventually this will grow to include a total of 18 antennas. The current concept considers 9-hour shifts at each complex, with a half hour overlap at the start as the complex assumes control, and another half hour overlap at the end as control is passed off to the next complex. Key enabling infrastructure includes network upgrades such that the Wide Area Network (WAN) connections between complexes is fast and reliable.

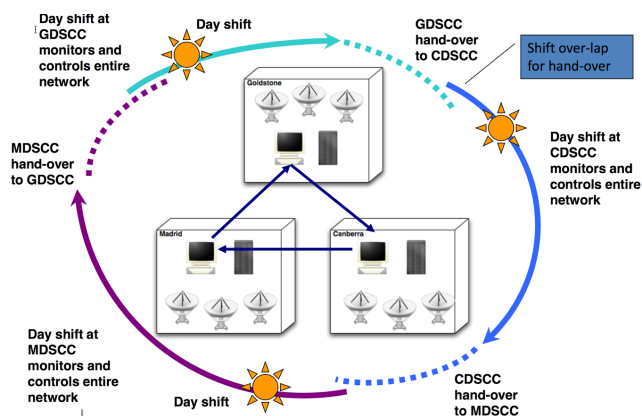


Fig 2 Diagram of Follow-the-Sun Operations showing controlling complex and handoff

From the point of view of network monitoring and control infrastructure, accomplishing remote operations requires significant under-the-hood changes to flow monitor data, logs, and event and status information reliably between complexes, and provide highly available connectivity and control software connections. Ultimately, operating an antenna remotely will be fully equivalent to operating one locally. If this were the extent of FtSO, no other changes would be required. However, the second major aspect of FtSO is to *increase by 50%* the number of links managed simultaneously by a Link Control Operator (LCO), from two to three. This brings in an additional set of changes related to helping manage the LCO workload and efficiency: among these are the following:

- integrating scheduling, service preparation, and real-time asset status to rapidly respond to changes
- link complexity scheduling and management, to proactively avoid work overload and to suggest efficient activity groupings
- a “self service portal” for missions to make realtime or near-realtime changes to their DSN activities, without requiring manual intervention by LCOs
- complex event processing — to process and aggregate current, historical, and planned information to provide LCOs with unprecedented contextual insight and data mining capabilities

In the following we cover these topics in more detail.

## 3. Responding to Change

One area that has been identified for improvement is that of rescheduling in response to changes, driven from various sources:

- users may have additional information or late changes to requirements for a variety of reasons
  - DSN assets (antennas, ground telecommunications equipment) may experience unexpected downtimes that require adjustments to the schedule to accommodate
  - spacecraft emergencies may occur that require extra tracking or changes to existing scheduled activities
- For many missions that are sequenced well in advance, late changes cannot be readily accommodated.

The DSN scheduling software is called Service Scheduling Software, or S<sup>3</sup>[4,5]. It was initially applied to the mid-range phase of the process, but is being extended to cover all three phases. S<sup>3</sup> provides a Javascript-based HTML5 web application and integrated database[10] through which users can directly enter their own scheduling requirements and verify their correctness before the submission deadline. The database in which requirements are stored is logically divided into “master” and “workspace” areas. There is a single master schedule representing mission-approved requirements and DSN activities (tracks). Each user can create an arbitrary number of workspace schedules, initially either empty or based on the contents of the master, within which they can conduct studies and ‘what if’ investigations, or keep a baseline for comparison with the master. These workspaces are by default private to the individual user, but can be shared as readable or read-write to any number of other users. Shared workspaces can be viewed and updated in realtime: while there can only be one writer at a time, any number of other users can view a workspace and see it automatically update as changes are made. These aspects of the web application architecture and database design support the collaborative and shared development nature of the DSN schedule.

In addition, S<sup>3</sup> offers specialized features to facilitate collaboration, including an integrated wiki for annotated discussion of negotiation proposals, integrated chat, notifications of various events, and a propose/concur/reject/counter workflow manager to support change proposals. Details on the design and use of the S<sup>3</sup> collaboration features[6] and the scheduling engine[7,8] are provided elsewhere.

In the FtSO era, S<sup>3</sup> will be made available for use at the complexes directly, to enter and manage activities such as maintenance and engineering. In addition, operators will be able to make other schedule changes to help manage their workload, such as starting setup for an activity earlier than usual, or extending teardown later. Such changes do not impact mission users of the DSN, but give the operations staff more flexibility.

In addition, with the integration of S<sup>3</sup> to support real-time, and access to realtime asset status as well as detailed requirements and flexibilities of individual activities, S<sup>3</sup> can be used to generate alternative rescheduling options when late breaking schedule changes occur. These can be due to any of the reasons noted above that can affect the real-time schedule. In today’s operations, such changes require a great deal of back and forth between the users and operations staff to come up with minimal impact schedule changes. Use of automated scheduling s/w to provide suggestions and options is expected to help facilitate this exchange.

Fig. 3 shows an example of this concept, representing a schedule change scenario in response to extended downtime on an antenna due to a problem discovered during routine maintenance. Instead of maintenance ending as planned, it is extended by 12 hours and thus impacts several upcoming mission activities on the affected antenna.

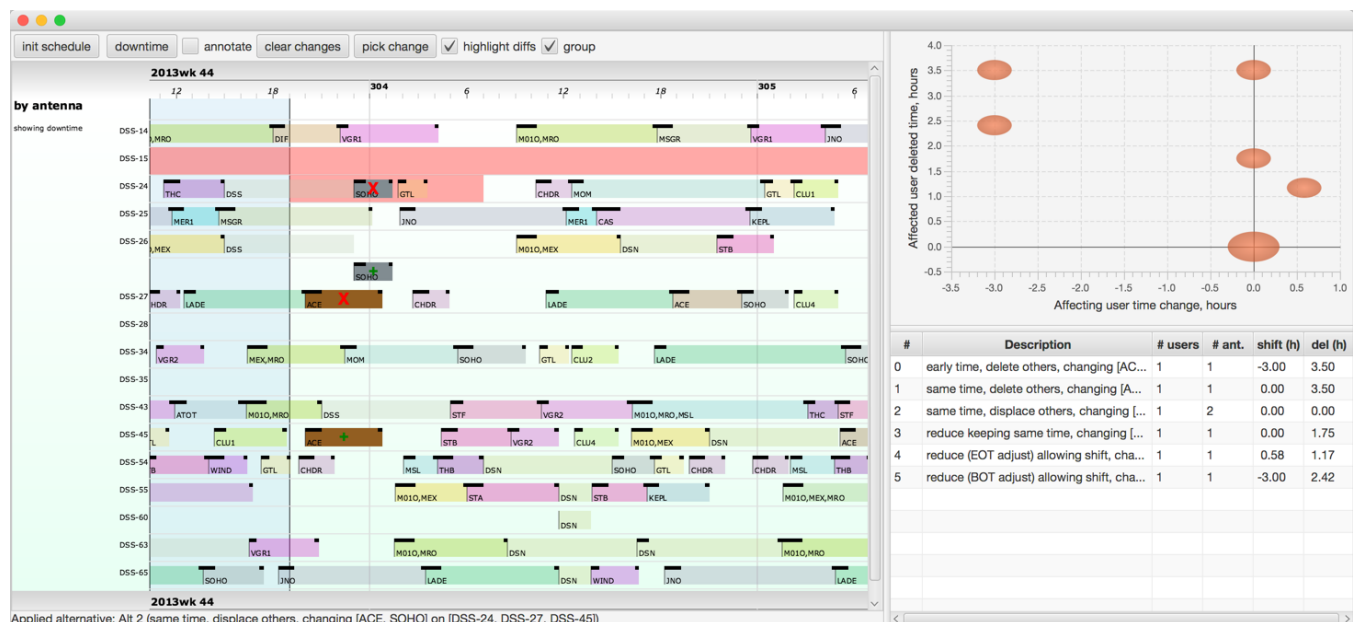


Fig 3. Example of rescheduling in response to asset availability change. A portfolio of options is illustrated and scored on the right.

By invoking a series of rescheduling strategies to fix the problem and score alternative solutions, the end user can be presented with options from which to choose. Since the activities are linked to the original requirements provided by the mission, such options are more likely to reflect real constraints and flexibilities, and minimize phone calls and delays to gather the equivalent information.

In this case, the strategy portfolio includes:

- fix the time (start, end) of the specified user and consider moving to different antennas and deleting activities that are in the way
- same, but consider shrinking activities that are in the way by the minimum needed to keep the specified user fixed in time on the different antenna
- allow the specified user to shift early or late, and consider reducing impacted activities from the start or end accordingly
- consider successive shifts to different antennas, keeping all start/end times the same

The two objective functions displayed on the Fig 3 XY bubble charts (top right) are (x) the time shift of the user being examined, and (y) how much time is lost from other impacted users. The size of the bubble indicates the number of antennas affected by the proposed change. All of these are criteria to be minimized, but this can rarely be accomplished simultaneously. Hence a tradeoff is required, and the GUI in Fig. 3 is intended to rapidly explore the trade space and assess potential recovery scenarios as quickly as possible.

In the illustrated situation, a solution is proposed that meets all user requirements but requires moving one user from DSS-27 to DSS-45 (a cross-complex shift that will be facilitated in remote operations), while the impacted user moves from DSS-24 to DSS-27. This preserves all the times and durations, but does require changes to two antennas.

## 4. Complexity Management

Not all activities are equally demanding, and when LCO are managing multiple activities at once it is easy to see that inadvertent overloading of the operations staff is a potential risk. As a result, we are investigating how to model the *complexity* of individual activities, and then to avoid overloading individual LCOs with too much work at one time. There are two major parts to this effort:

- a) *during schedule generation (weeks to months ahead of execution)*: to predict the occurrence of ‘spikes’ in loading and provide feedback to users so they can make adjustments early in the process before the schedule is firm; in addition, higher periods of link complexity could serve as early warning that additional or overtime staffing may be required to cover a particular time frame or critical event

- b) *during shift planning (hours to days ahead of execution)*: to determine a good assignment of work to operators that does not exceed threshold values for number of links or overall link complexity, and which, as much as possible, evenly distributes the work across the available operations staff (and also reduces the workload of the operations supervisors)

The concept of *link complexity* is intended to capture a measure of the workload of the Link Control Operator (LCO) while managing the three stages of a typical DSN-to-spacecraft link: setup, in-track, and teardown. This concept, to date, is not quantitative: there is no direct measure of how much concentration or mental energy is expended on a particular type of link. Several internal studies of link complexity have been conducted, which have investigated indicators of complexity based on measured quantities. One of the most promising indicators is termed Workload Intensity (WI), which reflects the rate at which the LCO issues commands or deals with interactive displays. High WI levels would be expected to correspond to high levels of required LCO attention and thus higher complexity.

Figure 4 depicts the measured 5-minute workload intensity relative to the beginning of each service provisioning phase in the DSN during a typical period of operations. Figure 4(a) shows the workload breakdown by type (directive or interactive display). Figure 4(b) shows the breakdown by spacecraft. A visual inspection clearly shows the following:


- the early minutes of each phase have the highest workload intensity for the setup and teardown phases
- a significant percentage of the workload is due to interactive displays
- workload intensity can vary significantly from spacecraft to spacecraft
- workload intensity tends to be more uniform during the actual service provisioning phase

A complexity estimation model or algorithm will need to account for the actual spacecraft being supported as well as the non-uniform distribution of workload during the time interval allocated to support the spacecraft.

In addition to the intrinsic complexity of the activity, the other major factor that needs to be taken into account in modeling complexity is that of external events, most notably shift change and handover. In the FtSO paradigm, each complex hands off ongoing activities to another complex when their day shift ends. During handoff, each LCO will be informing their successor of the state of the link and of any special considerations. During this time, the source and destination LCOs are more than normally occupied with their work, and so their capacity to take on new high-complexity activities is reduced.

As part of assessing the impact of higher link complexity in FtSO, and mitigating the risk of remote operations and 3LPO, we have developed a prototype and testbed for exploring link complexity models and scheduling algorithms. Our approach consists of a model of the operators

a)



b)




Figure 4: workload distribution across DSN activities by directive type (a) and mission (b), for each of the three service phases (setup, in-track, and teardown). “Workload Measure” is a count of operator-issued directives plus interactions with pop-up dialogs. The relative size of the squares is proportional to the number of service instances observed during the measured time period.

and a schedule of links as inputs. The operator model contains of two timelines:

- *Link count* – an integer resource measuring the number of links assigned to the operator.
- *Complexity value* – a floating point resource measuring the total complexity of the links assigned to the operator.

We have explored various assignment algorithms, as well as a GUI that allows for manual assignments and changes. Feedback from LCOs on the prototype has been very positive, and it will be developed further as part of the initial remote operations deployment. Details of the approach are provided in [9].

## 5. Self Service Portal

One of the areas that is known to take LCO attention during many passes is that of changes introduced by the mission via the voice line. Missions can change downlink data rates, turn command modulation on or off, and make a number of other unplanned changes, while remaining consistent with the overall parameters of the activity and its reserved assets. While this is not infrequent, it is not known the extent to which it limits operator attention, and thus represents a degree of risk to the goal of 3 links per operator. As a result, we are developing an approach to allow the mission user to make a specific set of allowed changes directly through a secure web portal, called the “self-service portal”.

Some of the factors that have to be considered in designing the portal are:

- *security* — because changes issued through the portal affect the spacecraft’s ground system configuration, it must be strictly controlled so that only an authorized user is involved, and that the changes are valid and safe
- *closed loop* — the user issuing requests through the portal needs clear visibility into what the current state of the spacecraft and ground system is, what the planned configuration will be, what changes are legal to request, and then that a request has been accepted and processed successfully. This is significantly more complicated than ‘fire and forget’ issuance of commands
- *integration with mission data systems* — missions already have existing ground system tools and GUIs and may need to integrate existing systems with the new DSN capability. This will require a very flexible interface, including web services as well as potential GUI-base portal functionality

DSN has been evolving the command sequencing interface between missions and the DSN towards the CCSDS standard for service management[3], and has developed a variant of the service profile and sequence specification designated “0211”. The 0211 specification is oriented towards the mission describing the spacecraft state as a function of time, along with the DSN services desired (telemetry, ranging, command, etc.). DSN then interprets this to provide

the correct assets, and asset configurations, at the correct time. This opens up the possibility to use the 0211 mission inputs to identify valid changes that could be made during a pass in realtime.

While there remains a great deal of design work before a DSN self-service portal could become a reality, it offers the potential to offload a common source of interruptions from the LCO’s direct involvement.

## 6. Complex Event Processing

“Complex Event Processing” (CEP) refers to analyzing data from the DSN assets as well as their components and systems, and correlating them with planned and historical data. In this case “Complex” refers to non-trivial — not to a given DSN facility. CEP is widely used in commercial settings and has grown in popularity in recent years[10]. In the DSN there are several potential areas of application as the Follow-the-Sun paradigm is put into place:

- *to detect anomalies, including deviations from predicted or historical behavior:* CEP rules can be defined to check complex conditions, through the use of time windowing, to look for events or trends that should be brought to the operator’s attention; conversely, items that might be blindly flagged by simple limit checking can be determined to be not of concern, given sufficient context and appropriate rules. CEP will naturally lend itself to a growing rulebase that can process and mine incoming data from heterogeneous sources. Among these are the predicted state (or state envelope) for comparison with actual data. CEP can also be configured to query archives for historical data and thus provide context for potential anomalies. In this usage mode, CEP can raise alarms that can be conveyed to users via GUIs or other mechanisms.
- *to diagnose, and recover from anomalies:* a more complicated and longer term goal is to have CEP diagnose problems, then recommend, or even carry out, recovery and mitigation steps. For example, a failing component might first be identified by its signature, then recommended to be sent to maintenance, and finally have a backup unit swapped in.
- *to filter, make consistent, and aggregate realtime data:* another application of CEP is to take existing realtime data streams and (1) support realtime transformations to reduce the naming and content discrepancies between logically related data, and (2) lower the cost to create customized aggregations for disparate DSN realtime use cases. For example, CEP can serve as middleware to present antenna-type specific information, such as pointing angles or weather data, in consistently named standards even when hardware may publish realtime data in non-standard form. Conversely, the same capability can also support customized realtime streams for niche use cases.

Because CEP can be used to aggregate and post-process data, in a highly scalable way, it offers a potential driver



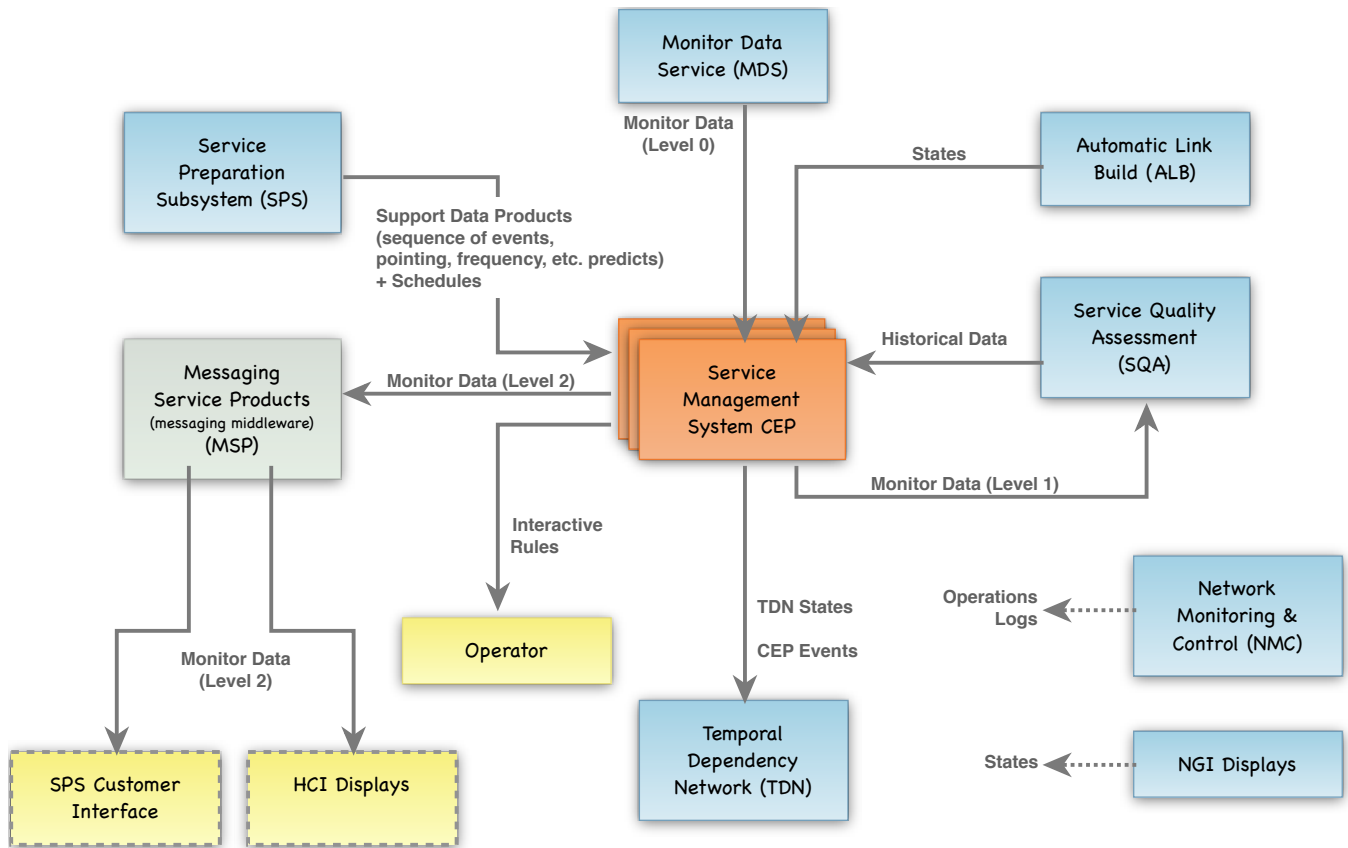


Figure 5: Context diagram of CEP showing interfaces to other DSN systems.

for new modes of human-computer interaction. New paradigms are being investigated that can give operators faster and clearer insight into the state of the network and how it is performing, as well as rapid drill-down into details behind deviations from expected behavior.

DSN ground system configurations, in support of spacecraft communication, vary depending on services requested, the front-end (antenna) selected, and the hardware assigned to the support. In addition, software configuration of the hardware allows for a large variety of spacecrafts with a variety of performance characteristics to be supported. Today, highly skilled operators monitor and oversee each of these spacecraft communications sessions. Spacecraft link budgets and thus telecom performance have significant variations due to spacecraft orientation and state, as well as earth-side environment — temperature, humidity, windspeed. With ~45 spacecrafts and 13 operating antennas, the number of potential configurations is enormous.

With Follow-the-Sun and 3LPO operations, individual operators will have to be cognizant of multiple simultaneous spacecraft tracks. Software assist is required. Temporal Dependency Network (TDN) scripts have been developed to assist operators with routine operations in running single and 2 LPO tracks. These scripts take as input Sequence of Event (SOE) predicts, and execute and monitor the systems from early setup, through pass execution, and into tear-

down. As we push the level of automation to higher levels, additional infrastructure is required.

The CEP infrastructure provides a mechanism to augment DSN's existing capabilities. Data from antenna pointing, frequency, round trip light time predicts, configuration tables, and SOEs are fed into a CEP engine. Streams of realtime data including several hundred monitor data streams and logs are also fed in. Algorithmic rules are run on the data streams and events are output through streams of middleware messages, logs, and TCP streams. Client applications can be connected to these output streams. Operators can then be alerted to conditions that can be watched over by the CEP system.

Figure 5 shows the CEP external interfaces. Monitor Data flows from the hardware systems to the Monitor Data Service (MDS), and on to CEP. Predicts flow from Service Preparation Subsystem (SPS) to CEP. TDNs interface bidirectionally — exchanging state information and subscribing to CEP events. The Automated Link Build (ALB) assembly interfaces with CEP to provide link build information. The Service Quality Assessment (SQA) assembly maintains a historical archive and provides the mechanism that allows CEP to have access to similar tracks that have occurred in the recent past.

Figure 6 shows a sample web interface to CEP. Realtime streaming data is organized and then output to a subscriber.

Dashboard	
CEP Data	
Raw Data	

Complex Event Processing Spark Dashboard	
Show 10 2 entries	Search:
Description	Value
AZ Actual MonData/CHOR65/AP65/AZANG	169.401026 degrees
AZ Actual MonData/LPO54/AP54/AzimuthAngle	169.261505 degrees
AZ Actual MonData/MOM34/AP34/AzimuthAngle	277.001801 degrees
AZ Actual MonData/nul55/AP55/AzimuthAngle	192.242645 degrees
AZ Actual MonData/VGR163/AP63/AZANG	191.014137 degrees
Downlink (physical) data channel number MonData/CHOR65/DOC2/DocChannel	2
Downlink (physical) data channel number MonData/GTL34/DOC7/DocChannel	7

Figure 6: Sample web interface to CEP.

Monitor data items are hashed, and a simple keyword search interface allows a user to view any streaming data in realtime. A client API is also available to allow this data to be accessed programmatically.

There are several use cases which particularly lend themselves to CEP techniques:

- Alert operations staff when spacecraft communications are not going according to plan: monitor tracking progress and compare to the plan. Use historical data to set thresholds. This allows for detection of false positives.
- Detect known “Master Discrepancies” when in progress: DSN maintains a list of known anomalies, and CEP can run rules that watch out for these by determining the signature of the current pass and correlating it with the signatures of these known anomalies. Then, operations can be alerted if a match is detected.
- User defined functions: some operators would like to have the system monitor particular variables and for them to be alerted in the event that the monitor detects some specific condition. This can alleviate the operator from having to monitor this on his or her own. Many such functions can be submitted to the system. These rules can evolve. If a subsystem delivers a new version, rules can be added to the CEP rulebase to watch for deviations from the norm.
- Global Realtime Status: CEP rules can monitor all tracks in progress and can output global combined status for the entire DSN.

## 7. Conclusions

The DSN Follow-the-Sun paradigm change represents a major change in the way the network is operated. Among the benefits will be lower operations costs, greater resilience, and the ability to support a larger network and user base. Many challenges have been identified in migrating the network from it’s current state into a remote operations configuration. The DSN passed its 50 year anniversary recently, and many parts of the network have legacy roots or residuals that complicate efforts at modernization. However, the initiatives described in this paper are all steps in the direction of improved cost-effective operations.

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. We gratefully acknowledge the support of the JPL Deep Space Network Project, and the Service Management System development and system engineering teams.

## Bibliography

- [1] J.B. Berner, J. Statman, Increasing the cost-efficiency of the DSN, in SpaceOps 2008: Heidelberg, Germany, 2008.
- [2] W.A. Imbriale, Large Antennas of the Deep Space Network, Wiley, 2003.
- [3] Space Communication Cross Support - Service Management - Service Specification, public.ccsds.org. (2009).
- [4] M.D. Johnston, D. Tran, Automated Scheduling for NASA's Deep Space Network, in: IWPSS 2008, Darmstadt, Germany, 2011.
- [5] M.D. Johnston, D. Tran, B. Arroyo, S. Sorensen, P. Tay, J. Carruth, et al., Automating Mid- and Long-Range Scheduling for NASA's Deep Space Network, in: SpaceOps 2012, Stockholm, Sweden, 2012.
- [6] J. Carruth, M.D. Johnston, A. Coffman, M. Wallace, B. Arroyo, S. Malhotra, A Collaborative Scheduling Environment for NASA's Deep Space Network, in: SpaceOps 2010, AIAA, Huntsville, AL, 2010.
- [7] M.D. Johnston, D. Tran, B. Arroyo, C. Page, Request-Driven Scheduling for NASA's Deep Space Network, in: Pasadena, CA, 2009.
- [8] S. Chien, M.D. Johnston, J. Frank, M. Giuliano, A. Kavelaars, C. Lenzen, et al., A generalized timeline representation, services, and interface for automating space mission operations, in: Proceedings of SpaceOps 2012, Stockholm, 2012.
- [9] D. Tran, M.D. Johnston, Automated Operator Link Assignment Scheduling for NASA's Deep Space Network, in: IWPSS 2015, Buenos Aires, Argentina, 2015.
- [10] D.C. Luckham, Event Processing for Business: Organizing the Real-Time Enterprise, John Wiley & Sons, Hoboken, 2012.